arXiv:0910.0651v2 [cs.IT] 21 Oct 2009

A Simpler Approach to Matrix Completion

Benjamin Recht

Department of Computer Sciences, University of Wisconsin-Madison
1210 W Dayton St, Madison, WI 53706
email: brecht@cs.wisc.edu

October 2009

Abstract

This paper provides the best bounds to date on the number of randomly sampled entries required to reconstruct an unknown low rank matrix. These results improve on prior work by Candès and Recht [4], Candès and Tao [7], and Keshavan, Montanari, and Oh [18]. The reconstruction is accomplished by minimizing the nuclear norm, or sum of the singular values, of the hidden matrix subject to agreement with the provided entries. If the underlying matrix satisfies a certain incoherence condition, then the number of entries required is equal to a quadratic logarithmic factor times the number of parameters in the singular value decomposition. The proof of this assertion is short, self contained, and uses very elementary analysis. The novel techniques herein are based on recent work in quantum information theory.

Keywords. Matrix completion, low-rank matrices, convex optimization, nuclear norm minimization, random matrices, operator Chernoff bound, compressed sensing.

1 Introduction

Recovering a low rank matrix from a given subset of its entries is a recurring problem in collaborative filtering [25], dimensionality reduction [20, 28], and multi-class learning [2, 22]. While a variety of heuristics have been developed across many disciplines, the general problem of finding the lowest rank matrix satisfying equality constraints is NP-hard. All known algorithms which can compute the lowest rank solution for all instances require time at least exponential in the dimensions of the matrix in both theory and practice [9].

In sharp contrast to such worst case pessimism, Candès and Recht showed that most low rank matrices could be recovered from most sufficiently large sets of entries by computing the matrix of minimum *nuclear norm* that agreed with the provided entries [4], and furthermore the revealed set of entries could comprise a vanishing fraction of the entire matrix. The nuclear norm is equal to the sum of the singular values of a matrix and is the best convex lower bound of the rank function on the set of matrices whose singular values are all bounded by 1. The intuition behind this heuristic is that whereas the rank function counts the number of nonvanishing singular values, the nuclear norm sums their amplitude, much like how the ℓ_1 norm is a useful surrogate for counting the number of nonzeros in a vector. Moreover, the nuclear norm can be minimized subject to equality constraints via semidefinite programming.

Nuclear norm minimization had long been observed to produce very low-rank solutions in practice (see, for example [3, 11, 12, 21, 26]), but only very recently was there any theoretical basis for when it produced the minimum rank solution. The first paper to provide such foundations was [24], where Recht, Fazel, and Parrilo developed probabilistic techniques to study average case behavior and showed that the nuclear norm heuristic could solve most instances of the rank minimization problem assuming the number of linear constraints was sufficiently large. The results in [24] inspired a groundswell of interest in theoretical guarantees for rank minimization, and these results lay the foundation for [4]. Candès and Recht's bounds were subsequently improved by Candès and Tao [7] and Keshavan, Montanari, and Oh [18] to show that one could, in special cases, reconstruct a low-rank matrix by observing a set of entries of size at most a polylogarithmic factor larger than the intrinsic dimension of the variety of rank r matrices.

This paper sharpens the results in [7, 18] to provide a bound on the number of entries required to reconstruct a low rank matrix which is optimal up to a small numerical constant and one logarithmic factor. The main theorem makes minimal assumptions about the low rank matrix of interest. Moreover, the proof is very short and relies on mostly elementary analysis.

In order to precisely state the main result, we need one definition. Candès and Recht observed that it is impossible to recover a matrix which is equal to zero in nearly all of its entries unless all of the entries of the matrix are observed (consider, for example, the rank one matrix which is equal to 1 in one entry and zeros everywhere else). In other words, the matrix cannot be mostly equal to zero on the observed entries. This motivated the following definition

Definition 1.1 *Let U be a subspace of $\mathbb{R}^n$ of dimension r and $\boldsymbol{P}_U$ be the orthogonal projection onto U. Then the* coherence *of U (vis-à-vis the standard basis $(\boldsymbol{e}_i)$) is defined to be*

$$\mu(U) \equiv \frac{n}{r} \max_{1 \leq i \leq n} \|\boldsymbol{P}_U \boldsymbol{e}_i\|^2. \tag{1.1}$$

Note that for any subspace, the smallest $\mu(U)$ can be is 1, achieved, for example, if U is spanned by vectors whose entries all have magnitude $1/\sqrt{n}$. The largest possible value for $\mu(U)$ is n/r which would correspond to any subspace that contains a standard basis element. If a matrix has row and column spaces with low coherence, then each entry can be expected to provide about the same amount of information.

Recall that the *nuclear norm* of an $n_1 \times n_2$ matrix $\boldsymbol{X}$ is the sum of the singular values of $\boldsymbol{X}$, $\|\boldsymbol{X}\|_* = \sum_{k=1}^{\min\{n_1,n_2\}} \sigma_k(\boldsymbol{X})$, where, here and below, $\sigma_k(\boldsymbol{X})$ denotes the kth largest singular value of $\boldsymbol{X}$. The main result of this paper is the following

Theorem 1.1 *Let $\boldsymbol{M}$ be an $n_1 \times n_2$ matrix of rank r with singular value decomposition $\boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^*$. Without loss of generality, impose the conventions $n_1 \leq n_2$, $\boldsymbol{\Sigma}$ is $r \times r$, $\boldsymbol{U}$ is $n_1 \times r$ and $\boldsymbol{V}$ is $n_2 \times r$. Assume that*

A0 *The row and column spaces have coherences bounded above by some positive μ_0.*

A1 *The matrix $\boldsymbol{U}\boldsymbol{V}^*$ has a maximum entry bounded by $\mu_1 \sqrt{r/(n_1 n_2)}$ in absolute value for some positive μ_1.*

Suppose m entries of $\boldsymbol{M}$ are observed with locations sampled uniformly at random. Then if

$$m \geq 32 \max\{\mu_1^2, \mu_0\}\, r(n_1 + n_2)\, \beta \log^2(2n_2) \tag{1.2}$$

for some $\beta > 1$, the minimizer to the problem

$$\begin{array}{ll} \textit{minimize} & \|\boldsymbol{X}\|_* \\ \textit{subject to} & X_{ij} = M_{ij} \quad (i,j) \in \Omega. \end{array} \tag{1.3}$$

is unique and equal to $\boldsymbol{M}$ with probability at least $1 - 6\log(n_2)(n_1+n_2)^{2-2\beta} - n_2^{2-2\beta^{1/2}}$.

The assumptions **A0** and **A1** were introduced in [4]. Both μ_0 and μ_1 may depend on r, n_1, or n_2. Moreover, note that $\mu_1 \leq \mu_0 \sqrt{r}$ by the Cauchy-Schwarz inequality. As shown in [4], both subspaces selected from the uniform distribution and spaces constructed as the span of singular vectors with bounded entries are not only incoherent with the standard basis, but also obey **A1** with high probability for values of μ_1 at most logarithmic in n_1 and/or n_2. Applying this theorem to the models studied in Section 2 of [4], we find that there is a numerical constant c_u such that $c_u r(n_1 + n_2)\log^5(n_2)$ entries are sufficient to reconstruct a rank r matrix whose row and column spaces are sampled from the Haar measure on the Grassmann manifold. If $r > \log(n_2)$, the number of entries can be reduced to $c_u r(n_1 + n_2)\log^4(n_2)$. Similarly, there is a numerical constant c_i such that $c_i \mu_0^2 r(n_1 + n_2)\log^3(n_2)$ entries are sufficient to recover a matrix of arbitrary rank r whose singular vectors have entries with magnitudes bounded by $\sqrt{\mu_0/n_1}$.

Theorem 1.1 greatly improves upon prior results. First of all, it has the weakest assumptions on the matrix to be recovered. In addition to assumption **A1**, Candès and Tao require a "strong incoherence condition" (see [7]) which is considerably more restrictive than the assumption **A0** in Theorem 1.1. Many of their results also require restrictions on the rank of $\boldsymbol{M}$, and their bounds depend superlinearly on μ_0. Keshavan *et al* require the matrix rank to be no more

than $\log(n_2)$, and require bounds on the maximum magnitude of the entries in $\boldsymbol{M}$ and the ratios $\sigma_1(\boldsymbol{M})/\sigma_r(\boldsymbol{M})$ and n_2/n_1. Theorem 1.1 makes no such assumptions about the rank, aspect ratio, nor condition number of $\boldsymbol{M}$. Moreover, (1.2) has a smaller log factor than [7], and features numerical constants that are both explicit and small.

Also note that there is not much room for improvement in the bound for m. It is a consequence of the coupon collector's problem that at least $n_2 \log n_2$ uniformly sampled entries are necessary just to guarantee that at least one entry in every row and column is observed with high probability. In addition, rank r matrices have $r(n_1 + n_2 - r)$ parameters, a fact that can be verified by counting the number of degrees of freedom in the singular value decomposition. Interestingly, Candès and Tao showed that $C\mu_0 n_2 r \log(n_2)$ entries were *necessary* for completion when the entries are sampled uniformly at random [7]. Hence, (1.2) is optimal up to a small numerical constant times $\log(n_2)$.

Most importantly, the proof of Theorem 1.1 is short and straightforward. Candès and Recht employed sophisticated tools from the study of random variables on Banach spaces including decoupling tools and powerful moment inequalities for the norms of random matrices. Candès and Tao rely on intricate moment calculations spanning over 30 pages. The present work only uses basic matrix analysis, elementary large deviation bounds, and a noncommutative version of Bernstein's Inequality proven here in the Appendix.

The proof of Theorem 1.1 is inspired by a recent paper in quanutm information which considered the problem of reconstructing the density matrix of a quantum ensemble using as few measurements as possible [16]. Their work adapted results from [4] and [5] to the quantum regime by using special algebraic properties of quantum measurements. Their proof followed a methodology analogous to the approach of Candès and Recht but had two main differences: they used a sampling with replacement model as a proxy for uniform sampling, and they deployed a powerful noncommutative Chernoff bound developed by Ahlswede and Winter for use in quantum information theory [1]. In this paper, I adapt these two strategies from [16] to the matrix completion problem. In section 3 I show how the sampling with replacement model bounds probabilities in the uniform sampling model, and present very short proofs of some of the main results in [4]. Surprisingly, this yields a simple proof of Theorem 1.1, provided in Section 4, which has the least restrictive assumptions of any assertion proven thus far.

2 Preliminaries and notation

Before continuing, let us survey the notations used throughout the paper. I closely follow the conventions established in [4], and invite the reader to consult this reference for a more thorough discussion of the matrix completion problem and the associated convex geometry. A thorough introduction to the necessary matrix analysis used in this paper can be found in [24].

Matrices are bold capital, vectors are bold lowercase and scalars or entries are not bold. For example, $\boldsymbol{X}$ is a matrix, and X_{ij} its (i, j)th entry. Likewise $\boldsymbol{x}$ is a vector, and x_i its ith component. If $\boldsymbol{u}_k \in \mathbb{R}^n$ for $1 \leq k \leq d$ is a collection of vectors, $[\boldsymbol{u}_1, \ldots, \boldsymbol{u}_d]$ will denote the $n \times d$ matrix whose kth column is $\boldsymbol{u}_k$. $\boldsymbol{e}_k$ will denote the kth standard basis vector in $\mathbb{R}^d$, equal to 1 in component k and 0 everywhere else. The dimension of $\boldsymbol{e}_k$ will always be clear from context. $\boldsymbol{X}^*$ and $\boldsymbol{x}^*$ denote the transpose of matrices $\boldsymbol{X}$ and vectors $\boldsymbol{x}$ respectively.

A variety of norms on matrices will be discussed. The spectral norm of a matrix is denoted by $\|\boldsymbol{X}\|$. The Euclidean inner product between two matrices is $\langle \boldsymbol{X}, \boldsymbol{Y} \rangle = \mathrm{Tr}(\boldsymbol{X}^* \boldsymbol{Y})$, and the corresponding Euclidean norm, called the Frobenius or Hilbert-Schmidt norm, is denoted $\|\boldsymbol{X}\|_F$. That is, $\|\boldsymbol{X}\|_F = \langle \boldsymbol{X}, \boldsymbol{X} \rangle^{1/2}$. The nuclear norm of a matrix $\boldsymbol{X}$ is $\|\boldsymbol{X}\|_*$. The maximum entry of $\boldsymbol{X}$ (in absolute value) is denoted by $\|\boldsymbol{X}\|_\infty \equiv \max_{ij} |X_{ij}|$. For vectors, the only norm applied is the usual Euclidean ℓ_2 norm, simply denoted as $\|\boldsymbol{x}\|$.

Linear transformations that act on matrices will be denoted by calligraphic letters. In particular, the identity operator will be denoted by $\mathcal{I}$. The spectral norm (the top singular value) of such an operator will be denoted by $\|\mathcal{A}\| = \sup_{\boldsymbol{X}:\|\boldsymbol{X}\|_F \leq 1} \|\mathcal{A}(\boldsymbol{X})\|_F$.

Fix once and for all a matrix $\boldsymbol{M}$ obeying the assumptions of Theorem 1.1. Let $\boldsymbol{u}_k$ (respectively $\boldsymbol{v}_k$) denote the kth column of $\boldsymbol{U}$ (respectively $\boldsymbol{V}$). Set $U \equiv \mathrm{span}\,(\boldsymbol{u}_1, \ldots, \boldsymbol{u}_r)$, and $V \equiv \mathrm{span}\,(\boldsymbol{v}_1, \ldots, \boldsymbol{v}_r)$. Also assume, without loss of generality, that $n_1 \leq n_2$. It is convenient to introduce the orthogonal decomposition $\mathbb{R}^{n_1 \times n_2} = T \oplus T^\perp$ where T is the linear space spanned by elements of the form $\boldsymbol{u}_k \boldsymbol{y}^*$ and $\boldsymbol{x} \boldsymbol{v}_k^*$, $1 \leq k \leq r$, where $\boldsymbol{x}$ and $\boldsymbol{y}$ are arbitrary, and $T^\perp$ is its orthogonal complement. $T^\perp$ is the subspace of matrices spanned by the family $(\boldsymbol{x}\boldsymbol{y}^*)$, where $\boldsymbol{x}$ (respectively $\boldsymbol{y}$) is any vector orthogonal to U (respectively V).

The orthogonal projection $\mathcal{P}_T$ onto T is given by

$$\mathcal{P}_T(\boldsymbol{Z}) = \boldsymbol{P}_U \boldsymbol{Z} + \boldsymbol{Z}\boldsymbol{P}_V - \boldsymbol{P}_U \boldsymbol{Z} \boldsymbol{P}_V, \tag{2.1}$$

where $\boldsymbol{P}_U$ and $\boldsymbol{P}_V$ are the orthogonal projections onto U and V respectively. Note here that while $\boldsymbol{P}_U$ and $\boldsymbol{P}_V$ are matrices, $\mathcal{P}_T$ is a linear operator mapping matrices to matrices. The orthogonal projection onto $T^\perp$ is given by

$$\mathcal{P}_{T^\perp}(\boldsymbol{Z}) = (\mathcal{I} - \mathcal{P}_T)(\boldsymbol{Z}) = (\boldsymbol{I}_{n_1} - \boldsymbol{P}_U)\boldsymbol{Z}(\boldsymbol{I}_{n_2} - \boldsymbol{P}_V)$$

where $\boldsymbol{I}_d$ denotes the $d \times d$ identity matrix. It follows from the definition (2.1) of $\mathcal{P}_T$ that

$$\mathcal{P}_T(\boldsymbol{e}_a \boldsymbol{e}_b^*) = (\boldsymbol{P}_U \boldsymbol{e}_a)\boldsymbol{e}_b^* + \boldsymbol{e}_a(\boldsymbol{P}_V \boldsymbol{e}_b)^* - (\boldsymbol{P}_U \boldsymbol{e}_a)(\boldsymbol{P}_V \boldsymbol{e}_b)^*.$$

This gives

$$\|\mathcal{P}_T(\boldsymbol{e}_a \boldsymbol{e}_b^*)\|_F^2 = \langle \mathcal{P}_T(\boldsymbol{e}_a \boldsymbol{e}_b^*), \boldsymbol{e}_a \boldsymbol{e}_b^* \rangle = \|\boldsymbol{P}_U \boldsymbol{e}_a\|^2 + \|\boldsymbol{P}_V \boldsymbol{e}_b\|^2 - \|\boldsymbol{P}_U \boldsymbol{e}_a\|^2 \|\boldsymbol{P}_V \boldsymbol{e}_b\|^2 .$$

Since $\|\boldsymbol{P}_U \boldsymbol{e}_a\|^2 \leq \mu(U) r / n_1$ and $\|\boldsymbol{P}_V \boldsymbol{e}_b\|^2 \leq \mu(V) r / n_2$,

$$\|\mathcal{P}_T(\boldsymbol{e}_a \boldsymbol{e}_b^*)\|_F^2 \leq \max\{\mu(U), \mu(V)\} r \frac{n_1 + n_2}{n_1 n_2} \leq \mu_0 r \frac{n_1 + n_2}{n_1 n_2} \tag{2.2}$$

I will make frequent use of this calculation throughout the sequel.

3 Sampling with Replacement

As discussed above, the main contribution of this work is an analysis of uniformly sampled sets of entries via the study of a sampling with replacement model. All of the previous work [4, 7, 18] studied a Bernoulli sampling model as a proxy for uniform sampling. There, each entry was revealed independently with probability equal to p. In all of these results, the theorem statements concerned sampling sets of m entries uniformly, but it was shown that probability of failure under Bernoulli sampling with $p = \frac{m}{n_1 n_2}$ closely approximated the probability of failure under uniform sampling. The present work will analyze the situation where each entry index is sampled independently from the uniform distribution on $\{1, \ldots, n_1\} \times \{1, \ldots, n_2\}$. This modification of the sampling model gives rise to all of the simplifications below.

It would appear that sampling with replacement is not suitable for analyzing matrix completion as one might encounter duplicate entries. However, just as is the case with Bernoulli sampling, bounding the likelihood of error when sampling with replacement allows us to bound the probability of the nuclear norm heuristic failing under uniform sampling.

Proposition 3.1 *The probability that the nuclear norm heuristic fails when the set of observed entries is sampled uniformly from the collection of sets of size m is less than or equal to the probability that the heuristic fails when m entries are sampled independently with replacement.*

Proof The proof follows the argument in Section II.C of [6]. Let Ω' be a collection of m entries, each sampled independently from the uniform distribution on $\{1, \ldots, n_1\} \times \{1, \ldots, n_2\}$. Let Ω_k denote a set of entries of size k sampled uniformly from all collections of entries of size k. It follows that

$$\begin{aligned}
\mathbb{P}(\text{Failure}(\Omega')) &= \sum_{k=0}^{m} P(\text{Failure}(\Omega') \mid |\Omega'| = k) P(|\Omega'| = k) \\
&= \sum_{k=0}^{m} P(\text{Failure}(\Omega_k)) P(|\Omega'| = k) \\
&\geq P(\text{Failure}(\Omega_m)) \sum_{k=0}^{m} P(|\Omega'| = k) = P(\text{Failure}(\Omega_m)) .
\end{aligned}$$

Where the inequality follows because $P(\text{Failure}(\Omega_m)) \geq P(\text{Failure}(\Omega_{m'}))$ if $m \leq m'$. That is, the probability decreases as the number of entries revealed is increased. ■

Surprisingly, changing the sampling model makes most of the theorems from [4] simple consequences of a non-commutative variant of Bernstein's Inequality.

Theorem 3.2 (Noncommutative Bernstein Inequality) *Let $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_L$ be independent zero-mean random matrices of dimension $d_1 \times d_2$. Suppose $\rho_k^2 = \max\{\|\mathbb{E}[\boldsymbol{X}_k \boldsymbol{X}_k^*]\|, \|\mathbb{E}[\boldsymbol{X}_k^* \boldsymbol{X}_k]\|\}$ and $\|\boldsymbol{X}_k\| \leq M$ almost surely for all k. Then for any $\tau > 0$,*

$$\mathbb{P}\left[\left\|\sum_{k=1}^{L} \boldsymbol{X}_k\right\| > \tau\right] \leq (d_1 + d_2) \exp\left(\frac{-\tau^2/2}{\sum_{k=1}^{L} \rho_k^2 + M\tau/3}\right).$$

Note that in the case that $d_1 = d_2 = 1$, this is precisely the two sided version of the standard Bernstein Inequality. When the $\boldsymbol{X}_k$ are diagonal, this bound is the same as applying the standard Bernstein Inequality and a union bound to the diagonal of the matrix summation. Furthermore, observe that the right hand side is less than $(d_1 + d_2)\exp(-\frac{3}{8}\tau^2/(\sum_{k=1}^{L} \rho_k^2))$ as long as $\tau \leq \frac{1}{M}\sum_{k=1}^{L} \rho_k^2$. This condensed form of the inequality will be used exclusively throughout. Theorem 3.2 is a corollary of an Chernoff bound for finite dimensional operators developed by Ahlswede and Winter [1]. A similar inequality for symmetric i.i.d. matrices is proposed in [16]. The proof is provided in the Appendix.

Let us now record two theorems, proven for the Bernoulli model in [4], that admit very simple proofs in the sampling with replacement model. The theorem statements requires some additional notation. Let $\Omega = \{(a_k, b_k)\}_{k=1}^{l}$ be a collection of indices sampled uniformly with replacement. Set $\mathcal{R}_\Omega$ to be the operator

$$\mathcal{R}_\Omega(\boldsymbol{Z}) = \sum_{k=1}^{|\Omega|} \langle \boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^*, \boldsymbol{Z} \rangle \boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^*.$$

Note that the (i,j)th component of $\mathcal{R}_\Omega(\boldsymbol{X})$ is zero unless $(i,j) \in \Omega$. For $(i,j) \in \Omega$, $\mathcal{R}_\Omega(\boldsymbol{X})$ is equal to X_{ij} times the multiplicity of $(i,j) \in \Omega$. Unlike in previous work on matrix completion, $\mathcal{R}_\Omega$ is not a projection operator if there are duplicates in Ω. Nonetheless, this does not adversely affect the argument, and $\mathcal{R}_\Omega(\boldsymbol{X}) = 0$ if and only if $X_{ab} = 0$ for all $(a,b) \in \Omega$. Moreover, we can show that the maximum duplication of any entry is always less than $\frac{8}{3}\log(n_2)$ with very high probability.

Proposition 3.3 *With probability at least $1 - n_2^{2-2\beta}$, the maximum number of repetitions of any entry in Ω is less than $\frac{8}{3}\beta\log(n_2)$ for $n_2 \geq 9$ and $\beta > 1$.*

Proof This assertion can be proven by applying a standard Chernoff bound for the Bernoulli distribution. Note that for a fixed entry, the probability it is sampled more than t times is equal to the probability of more than t heads occurring in a sequence of m tosses where the probability of a head is $\frac{1}{n_1 n_2}$. This probability can be upper bounded by

$$\mathbb{P}[\text{more than } t \text{ heads in } m \text{ trials}] \leq \left(\frac{m}{n_1 n_2 t}\right)^t \exp\left(t - \frac{m}{n_1 n_2}\right)$$

(see [17], for example). Applying the union bound over all of the $n_1 n_2$ entries and the fact that $\frac{m}{n_1 n_2} < 1$, we have

$$\mathbb{P}[\text{any entry is selected more than } \tfrac{8}{3}\beta\log(n_2) \text{ times}] \leq n_1 n_2 \left(\tfrac{8}{3}\beta\log(n_2)\right)^{-\frac{8}{3}\beta\log(n_2)} \exp\left(\tfrac{8}{3}\beta\log(n_2)\right) \leq n_2^{2-2\beta}$$

when $n_2 \geq 9$. ■

This application of the Chernoff bound is very crude, and much tighter bounds can be derived using more careful analysis. For example in [15], the maximum oversampling is shown to be bounded by $O(\frac{\log(n_2)}{\log\log(n_2)})$. For our purposes here, the loose upper bound provided by Proposition 3.3 will be more than sufficient.

In addition to this bound on the norm of $\mathcal{R}_\Omega$, the following theorem asserts that the operator $\mathcal{P}_T \mathcal{R}_\Omega \mathcal{P}_T$ is also very close to an isometry on T if the number of sampled entries is sufficiently large. This result is analgous to the Theorem 4.1 in [4] for the Bernoulli model, whose proof uses several powerful theorems from the study of probability in Banach spaces. Here, one only needs to compute a few low order moments and then apply Theorem 3.2.

Theorem 3.4 *Suppose Ω is a set of entries of size m sampled independently and uniformly with replacement. Then for all $\beta > 1$,*

$$\frac{n_1 n_2}{m} \left\| \mathcal{P}_T \mathcal{R}_\Omega \mathcal{P}_T - \frac{m}{n_1 n_2} \mathcal{P}_T \right\| \leq \sqrt{\frac{16 \mu_0 r (n_1 + n_2)\, \beta \log(n_2)}{3m}}$$

with probability at least $1 - 2n_2^{2-2\beta}$ provided that $m > \frac{16}{3} \mu_0 r (n_1 + n_2)\, \beta \log(n_2)$.

Proof Decompose any matrix $\boldsymbol{Z}$ as $\boldsymbol{Z} = \sum_{ab} \langle \boldsymbol{Z}, \boldsymbol{e}_a \boldsymbol{e}_b^* \rangle \boldsymbol{e}_a \boldsymbol{e}_b^*$ so that

$$\mathcal{P}_T(\boldsymbol{Z}) = \sum_{ab} \langle \mathcal{P}_T(\boldsymbol{Z}), \boldsymbol{e}_a \boldsymbol{e}_b^* \rangle \boldsymbol{e}_a \boldsymbol{e}_b^* = \sum_{ab} \langle \boldsymbol{Z}, \mathcal{P}_T(\boldsymbol{e}_a \boldsymbol{e}_b^*) \rangle \boldsymbol{e}_a \boldsymbol{e}_b^*. \tag{3.1}$$

For $k = 1, \ldots, m$ sample (a_k, b_k) from $\{1, \ldots, n_1\} \times \{1, \ldots, n_2\}$ uniformly with replacement. Then $\mathcal{R}_\Omega \mathcal{P}_T(\boldsymbol{Z}) = \sum_{k=1}^m \langle \boldsymbol{Z}, \mathcal{P}_T(\boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^*) \rangle \, \boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^*$ which gives

$$(\mathcal{P}_T \mathcal{R}_\Omega \mathcal{P}_T)(\boldsymbol{Z}) = \sum_{k=1}^m \langle \boldsymbol{Z}, \mathcal{P}_T(\boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^*) \rangle \, \mathcal{P}_T(\boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^*).$$

Now the fact that the operator $\mathcal{P}_T \mathcal{R}_\Omega \mathcal{P}_T$ does not deviate from its expected value

$$\mathbb{E}(\mathcal{P}_T \mathcal{R}_\Omega \mathcal{P}_T) = \mathcal{P}_T (\mathbb{E}\, \mathcal{R}_\Omega) \mathcal{P}_T = \mathcal{P}_T (\frac{m}{n_1 n_2} \mathcal{I}) \mathcal{P}_T = \frac{m}{n_1 n_2} \mathcal{P}_T$$

in the spectral norm can be proven using the Noncommutative Bernstein Inequality.

To proceed, define the operator $\mathcal{T}_{ab}$ which maps $\boldsymbol{Z}$ to $\langle \mathcal{P}_T(\boldsymbol{e}_a \boldsymbol{e}_b^*), \boldsymbol{Z} \rangle \mathcal{P}_T(\boldsymbol{e}_a \boldsymbol{e}_b^*)$. This operator is rank one, has operator norm $\|\mathcal{T}_{ab}\| = \|\mathcal{P}_T(\boldsymbol{e}_a \boldsymbol{e}_b^*)\|_F^2$, and we have $\mathcal{P}_T = \sum_{a,b} \mathcal{T}_{ab}$ by (3.1). Hence, for $k = 1, \ldots, m$, $\mathbb{E}[\mathcal{T}_{a_k b_k}] = \frac{1}{n_1 n_2} \mathcal{P}_T$.

Observe that if $\boldsymbol{A}$ and $\boldsymbol{B}$ are positive semidefinite, we have $\|\boldsymbol{A} - \boldsymbol{B}\| \leq \max\{\|\boldsymbol{A}\|, \|\boldsymbol{B}\|\}$. Using this fact, we can compute the bound

$$\|\mathcal{T}_{a_k b_k} - \tfrac{1}{n_1 n_2} \mathcal{P}_T\| \leq \max\{\|\mathcal{P}_T(\boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^*)\|_F^2, \tfrac{1}{n_1 n_2}\} \leq \mu_0 r \frac{n_1 + n_2}{n_1 n_2},$$

where the final inequality follows from (2.2). We also have

$$\begin{aligned}
\|\, \mathbb{E}[(\mathcal{T}_{a_k b_k} - \tfrac{1}{n_1 n_2} \mathcal{P}_T)^2]\| &= \|\, \mathbb{E}[\|\mathcal{P}_T(\boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^*)\|_F^2 \mathcal{T}_{a_k b_k}] - \frac{1}{n_1^2 n_2^2} \mathcal{P}_T]\| \\
&\leq \max\{\|\, \mathbb{E}[\|\mathcal{P}_T(\boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^*)\|_F^2 \mathcal{T}_{a_k b_k}]\|, \frac{1}{n_1^2 n_2^2}\} \\
&\leq \max\{\|\, \mathbb{E}[\mathcal{T}_{a_k b_k}]\| \, \mu_0 r \frac{n_1 + n_2}{n_1 n_2}, \frac{1}{n_1^2 n_2^2}\} \leq \mu_0 r \frac{n_1 + n_2}{n_1^2 n_2^2}
\end{aligned}$$

The theorem now follows by applying the Noncommutative Bernstein Inequality. ∎

The next theorem is an analog of Theorem 6.3 in [4] or Lemma 3.2 in [18]. This theorem asserts that for a fixed matrix, if one sets all of the entries not in Ω to zero it remains close to a multiple of the original matrix in the operator norm.

Theorem 3.5 *Suppose Ω is a set of entries of size m sampled independently and uniformly with replacement and let $\boldsymbol{Z}$ be a fixed $n_1 \times n_2$ matrix. Assume without loss of generality that $n_1 \leq n_2$, Then for all $\beta > 1$,*

$$\left\| \left(\frac{n_1 n_2}{m} \mathcal{R}_\Omega - \mathcal{I} \right) (\boldsymbol{Z}) \right\| \leq \sqrt{\frac{8 \beta n_1 n_2^2 \log(n_1 + n_2)}{3m}} \|\boldsymbol{Z}\|_\infty$$

with probability at least $1 - (n_1 + n_2)^{1-\beta}$ provided that $m > 6 \beta n_1 \log(n_1 + n_2)$.

Proof First observe that the operator norm can be upper bounded by a multiple of the matrix infinity norm

$$\|\boldsymbol{Z}\| = \sup_{\substack{\|\boldsymbol{x}\|=1 \\ \|\boldsymbol{y}\|=1}} \sum_{a,b} Z_{ab} y_a x_b \leq \left(\sum_{a,b} Z_{ab}^2 y_a^2 \right)^{1/2} \left(\sum_{a,b} x_b^2 \right)^{1/2} \leq \sqrt{n_2} \max_a \left(\sum_b Z_{ab}^2 \right)^{1/2} \leq \sqrt{n_1 n_2} \|\boldsymbol{Z}\|_\infty$$

Note that $\frac{n_1 n_2}{m} \mathcal{R}_\Omega(\boldsymbol{Z}) - \boldsymbol{Z} = \frac{1}{m} \sum_{k=1}^m n_1 n_2 Z_{a_k b_k} \boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^* - \boldsymbol{Z}$. This is a sum of zero-mean random matrices, and $\|n_1 n_2 Z_{a_k b_k} \boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^* - \boldsymbol{Z}\| \leq \|n_1 n_2 Z_{a_k b_k} \boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^*\| + \|\boldsymbol{Z}\| < \frac{3}{2} n_1 n_2 \|\boldsymbol{Z}\|_\infty$ for $n_1 \geq 2$. We also have

$$\begin{aligned} \left\| \mathbb{E}\left[(n_1 n_2 Z_{a_k b_k} \boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^* - \boldsymbol{Z})^* (n_1 n_2 Z_{a_k b_k} \boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^* - \boldsymbol{Z}) \right] \right\| &= \left\| n_1 n_2 \sum_{c,d} Z_{cd}^2 \boldsymbol{e}_d \boldsymbol{e}_d^* - \boldsymbol{Z}^* \boldsymbol{Z} \right\| \\ &\leq \max \left\{ \left\| n_1 n_2 \sum_{c,d} Z_{cd}^2 \boldsymbol{e}_d \boldsymbol{e}_d^* \right\|, \|\boldsymbol{Z}^* \boldsymbol{Z}\| \right\} \\ &\leq n_1 n_2^2 \|\boldsymbol{Z}\|_\infty^2 \end{aligned}$$

where we again use the fact that $\|\boldsymbol{A} - \boldsymbol{B}\| \leq \max\{\|\boldsymbol{A}\|, \|\boldsymbol{B}\|\}$ for positive semidefinite $\boldsymbol{A}$ and $\boldsymbol{B}$. A similar calculation holds for $(n_1 n_2 Z_{a_k b_k} \boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^* - \boldsymbol{Z})(n_1 n_2 Z_{a_k b_k} \boldsymbol{e}_{a_k} \boldsymbol{e}_{b_k}^* - \boldsymbol{Z})^*$. The theorem now follows by the Noncommutative Bernstein Inequality. ∎

Finally, the following Lemma is required to prove Theorem 1.1. Succinctly, it says that for a fixed matrix in T, the operator $\mathcal{P}_T \mathcal{R}_\Omega$ does not increase the matrix infinity norm.

Lemma 3.6 *Suppose Ω is a set of entries of size m sampled independently and uniformly with replacement and let $\boldsymbol{Z} \in T$ be a fixed $n_1 \times n_2$ matrix. Assume without loss of generality that $n_1 \leq n_2$. Then for all $\beta > 2$,*

$$\left\| \frac{n_1 n_2}{m} \mathcal{P}_T \mathcal{R}_\Omega(\boldsymbol{Z}) - \boldsymbol{Z} \right\|_\infty \leq \sqrt{\frac{8\beta \mu_0 r (n_1 + n_2) \log n_2}{3m}} \|\boldsymbol{Z}\|_\infty$$

with probability at least $1 - 2n_2^{2-\beta}$ provided that $m > \frac{8}{3} \beta \mu_0 r (n_1 + n_2) \log n_2$.

Proof This lemma can be proven using the standard Bernstein Inequality. For each matrix index (c, d), sample (a, b) uniformly at random to define the random variable $\xi_{cd} = \langle \boldsymbol{e}_c \boldsymbol{e}_d^*, n_1 n_2 \langle \boldsymbol{e}_a \boldsymbol{e}_b^*, \boldsymbol{Z} \rangle \mathcal{P}_T(\boldsymbol{e}_a \boldsymbol{e}_b^*) - \boldsymbol{Z} \rangle$. We have $\mathbb{E}[\xi_{cd}] = 0$, $|\xi_{cd}| \leq \mu_0 r (n_1 + n_2) \|\boldsymbol{Z}\|_\infty$, and

$$\begin{aligned} \mathbb{E}[\xi_{cd}^2] &= \frac{1}{n_1 n_2} \sum_{a,b} \langle \boldsymbol{e}_c \boldsymbol{e}_d^*, n_1 n_2 \langle \boldsymbol{e}_a \boldsymbol{e}_b^*, \boldsymbol{Z} \rangle \mathcal{P}_T(\boldsymbol{e}_a \boldsymbol{e}_b^*) - \boldsymbol{Z} \rangle^2 \\ &= n_1 n_2 \sum_{a,b} \langle \mathcal{P}_T(\boldsymbol{e}_c \boldsymbol{e}_d^*), \boldsymbol{e}_a \boldsymbol{e}_b^* \rangle^2 \langle \boldsymbol{e}_a \boldsymbol{e}_b^*, \boldsymbol{Z} \rangle^2 - Z_{cd}^2 \\ &\leq n_1 n_2 \|\mathcal{P}_T(\boldsymbol{e}_c \boldsymbol{e}_d^*)\|_F^2 \|\boldsymbol{Z}\|_\infty^2 \leq \mu_0 r (n_1 + n_2) \|\boldsymbol{Z}\|_\infty^2 . \end{aligned}$$

Since the (c, d) entry of $\frac{n_1 n_2}{m} \mathcal{P}_T \mathcal{R}_\Omega(\boldsymbol{Z}) - \boldsymbol{Z}$ is identically distributed to $\frac{1}{m} \sum_{k=1}^m \xi_{cd}^{(k)}$, where $\xi_{cd}^{(k)}$ are i.i.d. copies of ξ_{cd}, we have by Bernstein's Inequality and the union bound:

$$\Pr\left[\left\| \frac{n_1 n_2}{m} \mathcal{P}_T \mathcal{R}_\Omega(\boldsymbol{Z}) - \boldsymbol{Z} \right\|_\infty > \sqrt{\frac{8\beta \mu_0 r (n_1 + n_2) \log(n_2)}{3m}} \|\boldsymbol{Z}\|_\infty \right] \leq 2 n_1 n_2 \exp(-\beta \log(n_2)) \leq 2 n_2^{2-\beta} .$$

∎

4 Proof of Theorem 1.1

The proof follows the program developed in [16] which itself adapted the strategy proposed in [4]. The main idea is to approximate a dual feasible solution of (1.3) which certifies that $\boldsymbol{M}$ is the unique minimum nuclear norm solution. In [4] such a certificate was constructed via an infinite series using a construction developed in the compressed sensing literature [6, 13]. The terms in this series were then analyzed individually using the decoupling inequalities of de la Peña and Montgomery-Smith [10]. Truncating the infinite series after 4 terms gave their result. In [7], the authors bounded the contribution of $O(\log(n_2))$ terms in this series using intensive combinatorial analysis of each term. The insight in [16] was that, when sampling observations with replacement, a dual feasible solution could be closely approximated by a modified series where each term involved the product of independent random variables. This change in the sampling model allows one to avoid decoupling inequalities and gives rise to the dramatic simplification here.

To proceed, recall again that by Proposition 3.1 it suffices to consider the scenario when the entries are sampled independently and uniformly with replacement. I will first develop the main argument of the proof assuming many conditions hold with high probability. The proof is completed by subsequently bounding probability that all of these events hold. Suppose that

$$\frac{n_1 n_2}{m} \left\| \mathcal{P}_T \mathcal{R}_\Omega \mathcal{P}_T - \frac{m}{n_1 n_2} \mathcal{P}_T \right\| \leq \frac{1}{2}, \qquad \|\mathcal{R}_\Omega\| \leq \tfrac{8}{3} \beta^{1/2} \log(n_2). \tag{4.1}$$

Also suppose there exists a $\boldsymbol{Y}$ in the range of $\mathcal{R}_\Omega$ such that

$$\|\mathcal{P}_T(\boldsymbol{Y}) - \boldsymbol{U}\boldsymbol{V}^*\|_F \leq \sqrt{\frac{r}{2n_2}}, \qquad \|\mathcal{P}_{T^\perp}(\boldsymbol{Y})\| < \frac{1}{2} \tag{4.2}$$

If (4.1) holds, then for any $\boldsymbol{Z} \in \ker \mathcal{R}_\Omega$, $\mathcal{P}_T(\boldsymbol{Z})$ cannot be too large. Indeed, we have

$$0 = \|\mathcal{R}_\Omega(\boldsymbol{Z})\|_F \geq \|\mathcal{R}_\Omega \mathcal{P}_T(\boldsymbol{Z})\|_F - \|\mathcal{R}_\Omega \mathcal{P}_{T^\perp}(\boldsymbol{Z})\|_F .$$

Now observe that

$$\|\mathcal{R}_\Omega \mathcal{P}_T(\boldsymbol{Z})\|_F^2 = \langle \boldsymbol{Z}, \mathcal{P}_T \mathcal{R}_\Omega^2 \mathcal{P}_T(\boldsymbol{Z}) \rangle \geq \langle \boldsymbol{Z}, \mathcal{P}_T \mathcal{R}_\Omega \mathcal{P}_T(\boldsymbol{Z}) \rangle \geq \frac{m}{2n_1 n_2} \|\mathcal{P}_T(\boldsymbol{Z})\|_F^2$$

and $\|\mathcal{R}_\Omega \mathcal{P}_{T^\perp}(\boldsymbol{Z})\|_F \leq \frac{8}{3}\beta^{1/2}\log(n_2)\|\mathcal{P}_{T^\perp}(\boldsymbol{Z})\|_F$. Collecting these facts gives that for any $\boldsymbol{Z} \in \ker \mathcal{R}_\Omega$,

$$\|\mathcal{P}_{T^\perp}(\boldsymbol{Z})\|_F \geq \sqrt{\frac{9m}{128\beta n_1 n_2 \log^2(n_2)}} \|\mathcal{P}_T(\boldsymbol{Z})\|_F > \sqrt{\frac{2r}{n_2}} \|\mathcal{P}_T(\boldsymbol{Z})\|_F .$$

Now recall that $\|\boldsymbol{A}\|_* = \sup_{\|B\| \leq 1} \langle \boldsymbol{A}, \boldsymbol{B} \rangle$. For $\boldsymbol{Z} \in \ker \mathcal{R}_\Omega$, pick $\boldsymbol{U}_\perp$ and $\boldsymbol{V}_\perp$ such that $[\boldsymbol{U}, \boldsymbol{U}_\perp]$ and $[\boldsymbol{V}, \boldsymbol{V}_\perp]$ are unitary matrices and that $\langle \boldsymbol{U}_\perp \boldsymbol{V}_\perp^*, \mathcal{P}_{T^\perp}(\boldsymbol{Z}) \rangle = \|\mathcal{P}_{T^\perp}(\boldsymbol{Z})\|_*$. Then it follows that

$$\begin{aligned}
\|\boldsymbol{M} + \boldsymbol{Z}\|_* &\geq \langle \boldsymbol{U}\boldsymbol{V}^* + \boldsymbol{U}_\perp \boldsymbol{V}_\perp^*, \boldsymbol{M} + \boldsymbol{Z} \rangle \\
&= \|\boldsymbol{M}\|_* + \langle \boldsymbol{U}\boldsymbol{V}^* + \boldsymbol{U}_\perp \boldsymbol{V}_\perp^*, \boldsymbol{Z} \rangle \\
&= \|\boldsymbol{M}\|_* + \langle \boldsymbol{U}\boldsymbol{V}^* - \mathcal{P}_T(\boldsymbol{Y}), \mathcal{P}_T(\boldsymbol{Z}) \rangle + \langle \boldsymbol{U}_\perp \boldsymbol{V}_\perp^* - \mathcal{P}_{T^\perp}(\boldsymbol{Y}), \mathcal{P}_{T^\perp}(\boldsymbol{Z}) \rangle \\
&> \|\boldsymbol{M}\|_* - \sqrt{\frac{r}{2n_2}} \|\mathcal{P}_T(\boldsymbol{Z})\|_F + \frac{1}{2} \|\mathcal{P}_{T^\perp}(\boldsymbol{Z})\|_* \geq \|\boldsymbol{M}\|_* .
\end{aligned}$$

The first inequality holds from the variational characterization of the nuclear norm. We also used the fact that $\langle \boldsymbol{Y}, \boldsymbol{Z} \rangle = 0$ for all $\boldsymbol{Z} \in \ker \mathcal{R}_\Omega$. Thus, if a $\boldsymbol{Y}$ exists obeying (4.2), we have that for any $\boldsymbol{X}$ obeying $\mathcal{R}_\Omega(\boldsymbol{X} - \boldsymbol{M}) = \boldsymbol{0}$, $\|\boldsymbol{X}\|_* > \|\boldsymbol{M}\|_*$. That is, any if $\boldsymbol{X}$ has $M_{ab} = X_{ab}$ for all $(a, b) \in \Omega$, $\boldsymbol{X}$ has strictly larger nuclear norm than $\boldsymbol{M}$, and hence $\boldsymbol{M}$ is the unique minimizer of (1.3). The remainder of the proof shows that such a $\boldsymbol{Y}$ exists with high probability.

To this end, partition $1, \dots, m$ into p partitions of size q. By assumption, we may choose

$$q \geq \frac{128}{3} \max\{\mu_0, \mu_1^2\} r(n_1+n_2)\beta \log(n_1+n_2) \quad \text{and} \quad p \geq \frac{3}{4}\log(2n_2)\,.$$

Let Ω_j denote the set of indices corresponding to the jth partition. Note that each of these partitions are independent of one another when the indices are sampled with replacement. Assume that

$$\frac{n_1 n_2}{q}\left\|\mathcal{P}_T \mathcal{R}_{\Omega_k} \mathcal{P}_T - \frac{q}{n_1 n_2}\mathcal{P}_T\right\| \leq \frac{1}{2} \tag{4.3}$$

for all k. Define $\boldsymbol{W}_0 = \boldsymbol{U}\boldsymbol{V}^*$ and set $\boldsymbol{Y}_k = \frac{n_1 n_2}{q}\sum_{j=1}^{k} \mathcal{R}_{\Omega_j}(\boldsymbol{W}_{j-1})$, $\boldsymbol{W}_k = \boldsymbol{U}\boldsymbol{V}^* - \mathcal{P}_T(\boldsymbol{Y}_k)$ for $k = 1, \dots, p$. Then

$$\|\boldsymbol{W}_k\|_F = \left\|\boldsymbol{W}_{k-1} - \frac{n_1 n_2}{q}\mathcal{P}_T\mathcal{R}_{\Omega_k}(\boldsymbol{W}_{k-1})\right\|_F = \left\|(\mathcal{P}_T - \frac{n_1 n_2}{q}\mathcal{P}_T\mathcal{R}_{\Omega_k}\mathcal{P}_T)(\boldsymbol{W}_{k-1})\right\|_F \leq \frac{1}{2}\|\boldsymbol{W}_{k-1}\|_F\,,$$

and it follows that $\|\boldsymbol{W}_k\|_F \leq 2^{-k}\|\boldsymbol{W}_0\|_F = 2^{-k}\sqrt{r}$. Since $p \geq \frac{3}{4}\log(2n_2) \geq \frac{1}{2}\log_2(2n_2) = \log_2\sqrt{2n_2}$, then $\boldsymbol{Y} = \boldsymbol{Y}_p$ will satisfy the first inequality of (4.2). Also suppose that

$$\left\|\boldsymbol{W}_{k-1} - \frac{n_1 n_2}{q}\mathcal{P}_T\mathcal{R}_{\Omega_k}(\boldsymbol{W}_{k-1})\right\|_\infty \leq \frac{1}{2}\|\boldsymbol{W}_{k-1}\|_\infty \tag{4.4}$$

$$\left\|\left(\frac{n_1 n_2}{q}\mathcal{R}_{\Omega_j} - \mathcal{I}\right)(\boldsymbol{W}_{j-1})\right\| \leq \sqrt{\frac{8n_1 n_2^2 \beta \log n_2}{3q}}\|\boldsymbol{W}_{j-1}\|_\infty \tag{4.5}$$

for $k = 1, \dots, p$.

To see that $\|\mathcal{P}_{T^\perp}(\boldsymbol{Y}_p)\| \leq \frac{1}{2}$ when (4.4) and (4.5) hold, observe $\|\boldsymbol{W}_k\|_\infty \leq 2^{-k}\|\boldsymbol{U}\boldsymbol{V}^*\|_\infty$, and it follows that

$$\begin{aligned}
\|\mathcal{P}_{T^\perp}\boldsymbol{Y}_p\| &\leq \sum_{j=1}^{p} \|\tfrac{n_1 n_2}{q}\mathcal{P}_{T^\perp}\mathcal{R}_{\Omega_j}\boldsymbol{W}_{j-1}\| \\
&= \sum_{j=1}^{p} \|\mathcal{P}_{T^\perp}(\tfrac{n_1 n_2}{q}\mathcal{R}_{\Omega_j}\boldsymbol{W}_{j-1} - \boldsymbol{W}_{j-1})\| \\
&\leq \sum_{j=1}^{p} \|(\tfrac{n_1 n_2}{q}\mathcal{R}_{\Omega_j} - \mathcal{I})(\boldsymbol{W}_{j-1})\| \\
&\leq \sum_{j=1}^{p} \sqrt{\frac{8n_1 n_2^2\,\beta \log n_2}{3q}}\|\boldsymbol{W}_{j-1}\|_\infty \\
&= 2\sum_{j=1}^{p} 2^{-j}\sqrt{\frac{8n_1 n_2^2\,\beta \log n_2}{3q}}\|\boldsymbol{U}\boldsymbol{V}^*\|_\infty < \sqrt{\frac{32\mu_1^2 r n_2\,\beta \log n_2}{3q}} < 1/2
\end{aligned}$$

since $q > \frac{128}{3}\mu_1^2 r n_2 \beta \log(n_2)$. The first inequality follows from the triangle inequality. The second line follows because $\boldsymbol{W}_{j-1} \in T$ for all j. The third line follows because, for any $\boldsymbol{Z}$,

$$\|\mathcal{P}_{T^\perp}(\boldsymbol{Z})\| = \|(\boldsymbol{I}_{n_1} - \boldsymbol{P}_U)\boldsymbol{Z}(\boldsymbol{I}_{n_2} - \boldsymbol{P}_V)\| \leq \|\boldsymbol{Z}\|\,.$$

The fourth line applies (4.5). The next line follows from (4.4). The final line follows from the assumption **A1**.

All that remains is to bound the probability that all of the invoked events hold. With m satisfying the bound in the main theorem statement, the first inequality in (4.1) fails to hold with probability at most $2n_2^{2-2\beta}$ by Theorem 3.4, and the second inequality fails to hold with probability at most $n_2^{2-2\beta^{1/2}}$ by Proposition 3.3. For all k, (4.3) fails to hold with probability at most $2n_2^{2-2\beta}$, (4.4) fails to hold with probability at most $2n_2^{2-2\beta}$, and (4.5) fails to hold with probability at most $(n_1+n_2)^{1-2\beta}$. Summing these all together, all of the events hold with probability at least

$$1 - 6\log(n_2)(n_1+n_2)^{2-2\beta} - n_2^{2-2\beta^{1/2}}$$

by the union bound. This completes the proof.

5 Discussion and Conclusions

The results proven here are nearly optimal, but small improvements can possibly be made. The numerical constant 32 in the statement of the theorem may be reducible by more clever bookkeeping, and it may be possible to derive a linear dependence on the logarithm of the matrix dimensions. But further reduction is not possible because of the necessary conditions provided by Candès and Tao. One minor improvement that could be made would be to remove the assumption **A1**. For instance, while μ_1 is known to be small in most of the models of low rank matrices that have been analyzed, no one has shown that an assumption of the form **A1** is necessary for completion. Nonetheless, all prior results on matrix completion have imposed an assumption like **A1** [4, 7, 18], and it would be interesting to see if it can be removed as a requirement, or if it is somehow necessary.

Surprisingly, the simplicity of the argument presented here mostly arises from the abandonment of Bernoulli sampling in favor of sampling with replacement. It would be of interest to review results investigating noise robustness of matrix completion [5, 19] or deconvolution of sparse and low rank matrices [8] to see if results can be improved by appealing to sampling with replacement. Furthermore, since much of the work on rank minimization and matrix completion borrows tools from the compressed sensing community, it is of interest to revisit this related body of work and to see if proofs can be simplified or bounds can be improved there as well. The noncommutative versions of Chernoff and Bernstein' s Inequalities may be useful throughout machine learning and statistical signal processing, and a fruitful line of inquiry would examine how to apply these tools from quantum information to the study of classical signals and systems.

Acknowledgments

B.R. would like to thank Aram Harrow for introducing him to the operator Chernoff bound and many helpful clarifying conversations, Silvia Gandy for pointing out several typos in the original version of this manuscript, and Rob Nowak, Ali Rahimi, and Stephen Wright for many fruitful discussions about this paper.

References

[1] R. Ahlswede and A. Winter. Strong converse for identification via quantum channels. *IEEE Transactions on Information Theory*, 48(3):569–579, 2002.

[2] A. Argyriou, C. A. Micchelli, and M. Pontil. Convex multi-task feature learning. *Machine Learning*, 2008. Published online first at `http://www.springerlink.com/`.

[3] C. Beck and R. D'Andrea. Computational study and comparisons of LFT reducibility methods. In *Proceedings of the American Control Conference*, 1998.

[4] E. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 2008. To appear. Preprint available at `http://lanl.arxiv.org/abs/0805.4471`.

[5] E. J. Candès and Y. Plan. Matrix completion with noise. Submitted to *Proceedings of the IEEE*. Preprint available at `http://www-stat.stanford.edu/~candes/publications.html`, 2009.

[6] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52(2):489–509, 2006.

[7] E. J. Candès and T. Tao. The power of convex relaxation: Near-optimal matrix completion. Submitted for publication. Preprint available at `http://www-stat.stanford.edu/~candes/publications.html`, 2009.

[8] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky. Rank-sparsity incoherence for matrix decomposition. Submitted for publication. Preprint available at `http://ssg.mit.edu/group/venkatc/venkatc.shtml`, 2009.

[9] A. L. Chistov and D. Y. Grigoriev. Complexity of quantifier elimination in the theory of algebraically closed fields. In *Proceedings of the 11th Symposium on Mathematical Foundations of Computer Science*, volume 176 of *Lecture Notes in Computer Science*, pages 17–31. Springer Verlag, 1984.

[10] V. H. de la Peña and S. J. Montgomery-Smith. Decoupling inequalities for the tail probabilities of multivariate U-statistics. *Ann. Probab.*, 23(2):806–816, 1995.

[11] M. Fazel. *Matrix Rank Minimization with Applications*. PhD thesis, Stanford University, 2002.

[12] M. Fazel, H. Hindi, and S. Boyd. A rank minimization heuristic with application to minimum order system approximation. In *Proceedings of the American Control Conference*, 2001.

[13] J. J. Fuchs. On sparse representations in arbitrary redundant bases. *IEEE Transactions on Information Theory*, 50:1341–1344, 2004.

[14] S. Golden. Lower bounds for the Helmholtz function. *Physical Review*, 137B(4):B1127–1128, 1965.

[15] G. H. Gonnet. Expected length of the longest probe sequence in hash code searching. *Journal of the Association for Computing Machinery*, 28(2):289–304, 1981.

[16] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert. Quantum state tomography via compressed sensing. Preprint available at `http//arxiv.org/abs/0909.3304`, 2009.

[17] T. Hagerup and C. Rüb. A guided tour of chernoff bounds. *Information Processing Letters*, 33:305–308, 1990.

[18] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from a few entries. 2009. Preprint available at `http://arxiv.org/abs/0901.3150`.

[19] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. Preprint available at `http://arxiv.org/abs/0906.2027`, 2009.

[20] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15:215–245, 1995.

[21] M. Mesbahi and G. P. Papavassilopoulos. On the rank minimization problem over a positive semidefinite linear matrix inequality. *IEEE Transactions on Automatic Control*, 42(2):239–243, 1997.

[22] G. Obozinski, B. Taskar, and M. Jordan. Joint covariate selection and joint subspace selection for multiple classification problems. To Appear in *Journal of Statistics and Computing*. Preprint availble at `http://www.seas.upenn.edu/~taskar/`, 2009.

[23] D. Panchenko. MIT 18.465: Statistical Learning Theory. MIT Open Courseware `http://ocw.mit.edu/OcwWeb/Mathematics/18-465Spring-2007/CourseHome/`, 2007.

[24] B. Recht, M. Fazel, and P. Parrilo. Guaranteed minimum rank solutions of matrix equations via nuclear norm minimization. *SIAM Review*. To appear. Preprint Available at `http://pages.cs.wisc.edu/~brecht/publications.html`.

[25] J. D. M. Rennie and N. Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *Proceedings of the International Conference of Machine Learning*, 2005.

[26] N. Srebro. *Learning with Matrix Factorizations*. PhD thesis, Massachusetts Institute of Technology, 2004.

[27] C. J. Thompson. Inequality with applications in statistical mechanics. *Journal of Mathematical Physics*, 6(11):1812–1823, 1965.

[28] K. Q. Weinberger and L. K. Saul. Unsupervised learning of image manifolds by semidefinite programming. *International Journal of Computer Vision*, 70(1):77–90, 2006.

A Operator Chernoff Bounds

In this section, I present a proof of 3.2, and also provide new proofs of some probability bounds from quantum information theory. To review, a symmetric matrix $\boldsymbol{A}$ is positive semidefinite if all of its eigenvalues are nonnegative. If $\boldsymbol{A}$ and $\boldsymbol{B}$ are positive semidefinite matrices, $\boldsymbol{A} \preceq \boldsymbol{B}$ means $\boldsymbol{B} - \boldsymbol{A}$ is positive semidefinite. For square matrices $\boldsymbol{A}$, the matrix exponential will be denoted $\exp(\boldsymbol{A})$ and is given by the power series

$$\exp(\boldsymbol{A}) = \sum_{k=0}^{\infty} \frac{\boldsymbol{A}^k}{k!}$$

The following theorem is a generalization of Markov's inequality originally proven in [1]. My proof closely follows the standard proof of the traditional Markov inequality, and does not rely on discrete summations.

Theorem A.1 (Operator Markov Inequality [1]) *Let $\boldsymbol{X}$ be a random positive semidefinite matrix and $\boldsymbol{A}$ a fixed positive definite matrix. Then*

$$\mathbb{P}\left[\boldsymbol{X} \not\preceq \boldsymbol{A}\right] \leq \mathrm{Tr}(\mathbb{E}[\boldsymbol{X}]\boldsymbol{A}^{-1})$$

Proof Note that if $\boldsymbol{X} \not\preceq \boldsymbol{A}$, then $\boldsymbol{A}^{-1/2}\boldsymbol{X}\boldsymbol{A}^{-1/2} \not\preceq \boldsymbol{I}$, and hence $\|\boldsymbol{A}^{-1/2}\boldsymbol{X}\boldsymbol{A}^{-1/2}\| > 1$. Let $I_{\boldsymbol{X}\not\preceq\boldsymbol{A}}$ denote the indicator of the event $\boldsymbol{X} \not\preceq \boldsymbol{A}$. Then $I_{\boldsymbol{X}\not\preceq\boldsymbol{A}} \leq \mathrm{Tr}(\boldsymbol{A}^{-1/2}\boldsymbol{X}\boldsymbol{A}^{-1/2})$ as the right hand side is always nonnegative, and, if the left hand side equals 1, the trace of the right hand side must exceed the norm of the right hand side which is greater than 1. Thus we have

$$\mathbb{P}[\boldsymbol{X} \not\preceq \boldsymbol{A}] = \mathbb{E}[I_{\boldsymbol{X}\not\preceq\boldsymbol{A}}] \leq \mathbb{E}[\mathrm{Tr}(\boldsymbol{A}^{-1/2}\boldsymbol{X}\boldsymbol{A}^{-1/2})] = \mathrm{Tr}(\mathbb{E}[\boldsymbol{X}]\boldsymbol{A}^{-1}) \,.$$

where the last equality follows from the linearity and cyclic properties of the trace. ■

Next I will derive a noncommutative version of the Chernoff bound. This was also proven in [1] for i.i.d. matrices. The version stated here is more general in that the random matrices need not be identically distributed, but the proof is essentially the same.

Theorem A.2 (Noncommutative Chernoff Bound) *Let $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_n$ be independent symmetric random matrices in $\mathbb{R}^{d\times d}$. Let $\boldsymbol{A}$ be an arbitrary symmetric matrix. Then for any invertible $d \times d$ matrix $\boldsymbol{T}$*

$$\mathbb{P}\left[\sum_{k=1}^{n} \boldsymbol{X}_k \not\preceq n\boldsymbol{A}\right] \leq d \prod_{k=1}^{n} \|\mathbb{E}[\exp(\boldsymbol{T}\boldsymbol{X}_{\boldsymbol{k}}\boldsymbol{T}^* - \boldsymbol{T}\boldsymbol{A}\boldsymbol{T}^*)]\|$$

Proof The proof relies on an estimate from statistical physics which is stated here without proof.

Lemma A.3 (Golden-Thompson inequality [14, 27]) *For any symmetric matrices $\boldsymbol{A}$ and $\boldsymbol{B}$,*

$$\mathrm{Tr}(\exp(\boldsymbol{A}+\boldsymbol{B})) \leq \mathrm{Tr}((\exp \boldsymbol{A})(\exp \boldsymbol{B})) \,.$$

Much like the proof of the standard Chernoff bound, the theorem now follows from a long chain of inequalities.

$$\begin{aligned}
\mathbb{P}\left[\sum_{k=1}^{n} \boldsymbol{X}_k \not\preceq n\boldsymbol{A}\right] &= \mathbb{P}\left[\sum_{k=1}^{n} (\boldsymbol{X}_k - \boldsymbol{A}) \not\preceq 0\right] \\
&= \mathbb{P}\left[\sum_{k=1}^{n} \boldsymbol{T}(\boldsymbol{X}_k - \boldsymbol{A})\boldsymbol{T}^* \not\preceq 0\right] \\
&= \mathbb{P}\left[\exp\left(\sum_{k=1}^{n} \boldsymbol{T}(\boldsymbol{X}_k - \boldsymbol{A})\boldsymbol{T}^*\right) \not\preceq \boldsymbol{I}_d\right] \\
&\leq \mathrm{Tr}\left(\mathbb{E}\left[\exp\left(\sum_{k=1}^{n} \boldsymbol{T}(\boldsymbol{X}_k - \boldsymbol{A})\boldsymbol{T}^*\right)\right]\right) \\
&= \mathbb{E}\left[\mathrm{Tr}\left(\exp\left(\sum_{k=1}^{n} \boldsymbol{T}(\boldsymbol{X}_k - \boldsymbol{A})\boldsymbol{T}^*\right)\right)\right] \\
&\leq \mathbb{E}\left[\mathrm{Tr}\left(\exp\left(\sum_{k=1}^{n-1} \boldsymbol{T}(\boldsymbol{X}_k - \boldsymbol{A})\boldsymbol{T}^*\right)\exp\left(\boldsymbol{T}(\boldsymbol{X}_n - \boldsymbol{A})\boldsymbol{T}^*\right)\right)\right] \\
&\leq \mathbb{E}_{1,\ldots,n-1}\left[\mathrm{Tr}\left(\exp\left(\sum_{k=1}^{n-1} \boldsymbol{T}(\boldsymbol{X}_k - \boldsymbol{A})\boldsymbol{T}^*\right)\mathbb{E}[\exp\left(\boldsymbol{T}(\boldsymbol{X}_n - \boldsymbol{A})\boldsymbol{T}^*\right)]\right)\right] \\
&\leq \|\mathbb{E}[\exp\left(\boldsymbol{T}(\boldsymbol{X}_n - \boldsymbol{A})\boldsymbol{T}^*\right)]\|\,\mathbb{E}_{1,\ldots,n-1}\left[\mathrm{Tr}\left(\exp\left(\sum_{k=1}^{n-1} \boldsymbol{T}(\boldsymbol{X}_k - \boldsymbol{A})\boldsymbol{T}^*\right)\right)\right] \\
&\leq \prod_{k=2}^{n} \|\mathbb{E}[\exp\left(\boldsymbol{T}(\boldsymbol{X}_k - \boldsymbol{A})\boldsymbol{T}^*\right)]\|\,\mathbb{E}\left[\mathrm{Tr}\left(\exp\left(\boldsymbol{T}(\boldsymbol{X}_1 - \boldsymbol{A})\boldsymbol{T}^*\right)\right)\right] \\
&\leq d \prod_{k=1}^{n} \|\mathbb{E}[\exp\left(\boldsymbol{T}(\boldsymbol{X}_k - \boldsymbol{A})\boldsymbol{T}^*\right)]\|
\end{aligned}$$

Here, the first three lines follow from standard properties of the semidefinite ordering. The fourth line invokes the Operator Markov Inequality. The sixth line follows from the Golden-Thompson inequality. The seventh line follows from independence of the $\boldsymbol{X}_k$. The eighth line follows because for positive definite matrices $\mathrm{Tr}(\boldsymbol{AB}) \leq \mathrm{Tr}(\boldsymbol{A})\|\boldsymbol{B}\|$. This is just another statement of the duality between the nuclear and operator norms. The ninth line iteratively repeats the previous two steps. The final line follows because for a positive definite matrix $\boldsymbol{A}$, $\mathrm{Tr}(\boldsymbol{A})$ is the sum of the eigenvalues of $\boldsymbol{A}$, and all of the eigenvalues are at most $\|\boldsymbol{A}\|$. ■

Let us now turn to proving the Noncommutative Bernstein Inequality presented in Section 3. The authors in [16] proposed a similar inequality for symmetric i.i.d. random matrices with a slightly worse constant. The proof here is more general and follows the standard derivation of Bernstein's inequality.

Proof [of Theorem 3.2] Set

$$\boldsymbol{Y}_k = \begin{bmatrix} \mathbf{0} & \boldsymbol{X}_k \\ \boldsymbol{X}_k^* & \mathbf{0} \end{bmatrix}$$

Then $\boldsymbol{Y}_k$ are symmetric random variables, and for all k

$$\|\,\mathbb{E}[\boldsymbol{Y}_k^2]\| = \left\|\mathbb{E}\left[\begin{bmatrix} \boldsymbol{X}_k\boldsymbol{X}_k^* & \mathbf{0} \\ \mathbf{0} & \boldsymbol{X}_k^*\boldsymbol{X}_k \end{bmatrix}\right]\right\| = \max\{\|\,\mathbb{E}[\boldsymbol{X}_k\boldsymbol{X}_k^*]\|, \|\,\mathbb{E}[\boldsymbol{X}_k^*\boldsymbol{X}_k]\|\} = \rho_k^2\,.$$

Moreover, the maximum singular value of $\sum_{k=1}^L \boldsymbol{X}_k$ is equal to the maximum eigenvalue of $\sum_{k=1}^L \boldsymbol{Y}_k$. By Theorem A.2, we have for all $\lambda > 0$

$$\mathbb{P}\left[\left\|\sum_{k=1}^L \boldsymbol{X}_k\right\| > Lt\right] = \mathbb{P}\left[\sum_{k=1}^L \boldsymbol{Y}_k \not\preceq Lt\boldsymbol{I}\right] \leq (d_1+d_2)\exp(-L\lambda t)\prod_{k=1}^L \|\mathbb{E}[\exp(\lambda \boldsymbol{Y_k})]\|\,.$$

For each k, let $\boldsymbol{Y}_k = \boldsymbol{U}_k\boldsymbol{\Lambda}_k\boldsymbol{U}_k^*$ be an eigenvalue decomposition, where $\boldsymbol{\Lambda}_k$ is the diagonal matrix of the eigenvalues of $\boldsymbol{Y}_k$. In turn, it follows that for $s > 0$

$$-M^s\boldsymbol{Y}_k^2 \preceq -\boldsymbol{U}_kM^s\boldsymbol{\Lambda}_k^2\boldsymbol{U}_k^* \preceq \boldsymbol{U}_k\boldsymbol{\Lambda}_k^{2+s}\boldsymbol{U}_k^* = \boldsymbol{Y}_k^{2+s} \preceq \boldsymbol{U}_kM^s\boldsymbol{\Lambda}_k^2\boldsymbol{U}_k^* \preceq M^s\boldsymbol{Y}_k^2\,,$$

which then implies

$$\|\,\mathbb{E}[\boldsymbol{Y}_k^{s+2}]\| \leq M^s\|\,\mathbb{E}[\boldsymbol{Y}_k^2]\|\,. \tag{A.1}$$

For fixed k, we have

$$\begin{aligned}
\|\,\mathbb{E}[\exp(\lambda\boldsymbol{Y}_k)]\| &\leq \|\boldsymbol{I}\| + \sum_{j=2}^\infty \frac{\lambda^j}{j!}\|\,\mathbb{E}[\boldsymbol{Y}_k^j]\| \\
&\leq 1 + \sum_{j=2}^\infty \frac{\lambda^j}{j!}\|\,\mathbb{E}[\boldsymbol{Y}_k^2]\|M^{j-2} \\
&= 1 + \frac{\rho_k^2}{M^2}\sum_{j=2}^\infty \frac{\lambda^j}{j!}M^j = 1 + \frac{\rho_k^2}{M^2}(\exp(\lambda M) - 1 - \lambda M) \\
&\leq \exp\left(\frac{\rho_k^2}{M^2}(\exp(\lambda M) - 1 - \lambda M)\right)\,.
\end{aligned}$$

The first inequality follows from the triangle inequality and the fact that $\mathbb{E}[\boldsymbol{Y}_k] = 0$, the second inequality follows from (A.1), and the final inequality follows from the fact that $1 + x \leq \exp(x)$ for all x. Putting this together gives

$$\mathbb{P}\left[\left\|\sum_{k=1}^L \boldsymbol{X}_k\right\| > Lt\right] \leq (d_1+d_2)\exp\left(-\lambda Lt + \frac{\sum_{k=1}^L \rho_k^2}{M^2}(\exp(\lambda M) - 1 - \lambda M)\right)\,.$$

This final expression is now just a real number, and only has to be minimized as a function of λ. The theorem now follows by algebraic manipulation: the right hand side is minimized by setting $\lambda = \frac{1}{M}\log(1 + \frac{tLM}{\sum_{k=1}^L \rho_k^2})$, then basic approximations can be employed to complete the argument (see, for example [23], lectures 4 and 5). ■